

## **DETAILED ACTION**

### ***Response to Pre-Appeal Brief Request for Review***

The pre-appeal brief request for review filed 8 March 2010 has been entered. Claims 31-55 are pending. No claims are currently amended. Claims 1-30 are cancelled. No claims are new.

### ***Examiner's Amendment***

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Stephen Chang on 29 April 2010.

#### **Please amend the claims as follows:**

31. (Currently Amended) A method, comprising:  
creating a shadow cache configured to hold at least one modified block that has been physically redone and undone or that has been logically undone;  
collecting an active transaction list of all active transactions when a read-only transaction starts;  
determining a first log sequence number and a second log sequence number associated with the active transaction list, the first log sequence number and the second log sequence number specifying a set of transactions capable of being physically redone and undone or

logically undone, wherein the second log sequence number is greater than the first log sequence number;

loading a modified block into a read-only cache view, the modified block including a log sequence number;

comparing the log sequence number to both the first log sequence number and the second log sequence number;

performing physical redo operations or physical undo operations based upon the comparing as the another modified block is loaded into the read-only cache view, wherein the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the fist log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the modified block when the log sequence number is greater than the second log sequence number;

processing the active transaction list to physically undo an incomplete action and to logically undo an active transaction to generate transactional consistency in the read-only cache view;

creating one or more other modified blocks to generate the transactional consistency in the read-only cache view; and

determining the one or more other modified blocks being stored in the shadow cache and not an original database [[if]] when a shared cache overflows.

33. (Currently amended) The method of claim 31, wherein the original database is used [[if]] when the read-only cache view overflows a shared cache.

41. (Currently amended) The method of claim 40, further comprising using the back link log record to skip a portion of the transaction log that is irrelevant to undoing an uncommitted write transaction, wherein the back link log record is generated in the transaction log when there is an active read only transaction, [[if]] when the read-only transaction must be undone.

44. (Currently amended) A database system configured to restore a database to a consistent version supporting read-only uses, comprising:

a computer having a processor and memory;  
a log manager module configured to manage a transaction log of the database system;  
a cache manager module configured to manage a shared cache that stores one or more database blocks in a memory of the database system, and configured to create a write view of the database in the shared cache supporting read and write uses of the database and a read-only cache view of the database using the transaction log of the database, the read-only cache view being created in response to a read-only transaction of the database, the read-only cache view comprising the one or more database blocks of the shared cache that record a view of a version of the database at a time, wherein the cache manager is configured to use a non-transactional database to store the one or more database blocks that overflow the shared cache during use of the read-only cache view by the read-only transaction, wherein the steps to create the read-cache view include determining a first log sequence number and second log sequence number from the transaction log, wherein the second log sequence number is greater than the first log sequence number, loading a modified block including a log sequence number into the read-only cache

view, comparing the log sequence number to the first log sequence number and the second log sequence number, and performing physical redo operations or physical undo operations based upon the comparing as the modified block is loaded into the read-only cache view, wherein the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the first log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the modified block when the log sequence number is greater than the second log-sequence number; and

a transaction manager module configured to logically undo one or more transactions on demand, the one or more transactions having begun but having yet to commit upon starting the read-only transaction in order to construct the read-only cache view comprising a transactionally consistent prior version of the database, performing the read-only transaction using the read-only cache view without blocking performance of the one or more transactions involving a write operation using a write view of the database, and returning a result associated with the read-only transaction.

47. (Currently amended) The system of claim 46, wherein the non-transactional database is used [[if]] when the read-only cache view overflows the shared cache.

52. (Currently amended) The system of claim 51, wherein the cache manager is configured to traverse the shared cache looking for the one or more blocks to purge, and is further configured to purge the one or more blocks from the read-only cache view [[if]] when the

one or more blocks from the read-only cache view have been marked as closed when a new block allocation is requested in the shared cache.

***Reasons for Allowance***

The following is the examiner's statement of reasons for allowance:

The prior art of record, Klein et al., U.S. 6,631,374 ("Klein"), Ganesh et al., U.S. 6,192,377 ("Ganesh"), Weems, "Shadow Cache," University of Massachusetts, May 2004 ("Weems"), and Natarajan et al., "Log Sequence Numbers," University of Wisconsin, May 2003 ("Natarajan"), teach analogous art to the instant application, specifically, caching and logging transactions.

Klein teaches a method for selectively viewing temporal row data in a consistent read. Committed transactions effecting row data changes to a database are tracked and a stored system change number is assigned upon each committed transaction. Rollback segments for uncommitted transactions are stored. A queried selection of row data values from the database is performed as of a query time occurring prior to the commit time of at least one committed transaction. The ordered row data values contained in the rollback segments storing a transaction identifier for the at least one committed transaction are retrieved. See Abstract.

Ganesh teaches a method for determining whether a transaction can use a version of a data item. Transaction summary data specifies both a data buffer refresh time and a data buffer commit time. The data buffer refresh time is indicative of the latest time at which the version of data is guaranteed to be current. The data buffer commit time is indicative of the commit time of a most recently committed transaction from the one or more transactions that committed after

updating the version of the data item. A determination is made, based upon the transaction summary data and a snapshot time used by the first transaction, whether the version of the data item is to be provided to the first transaction. See Abstract.

Natarajan teaches that every log record is assigned a unique log sequence number.

Weems teaches that a shadow cache is a small separate cache storing address of each value evicted from the main cache.

In order to provide a consistent read, Klein teaches rolling back each active transaction to logically undo those transactions. Once all active transactions have been rolled back, if the transaction occurred earlier than the system change number corresponding to the requested query time, the entire transaction is also rolled back to logically undo the transaction. See Fig. 10.

Klein does not teach “performing physical redo operations or physical undo operations” based on “comparing the log sequence number to both the first log sequence number and the second log sequence number,” wherein “the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the fist log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the modified block when the log sequence number is greater than the second log sequence number,” as recited in independent claims 31 and 44.

Ganesh teaches determining whether a transaction can use a version of a data item by determining whether the snapshot time used by a transaction that requires access to a data block is less than the data buffer refresh time. If not, the transaction uses another version of the data

block. If, however, the snapshot time less than the data buffer refresh time, the method determines if the snapshot time is greater than the data buffer commit time. If not, the version of the data block in the data buffer is too new and the transaction are removed, for example by undoing or rolling back the changes. If, however, the snapshot time is greater than the data buffer commit time, then either changes are removed or the method uses that version of the data block, depending on whether the transaction is active or not. See Fig. 2. Ganesh does not teach “performing physical redo operations or physical undo operations” based on “comparing the log sequence number to both the first log sequence number and the second log sequence number,” wherein “the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the fist log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the modified block when the log sequence number is greater than the second log sequence number,” as recited in independent claims 31 and 44.

Natarajan teaches that every log record is assigned a unique log sequence number. Natarajan does not teach “performing physical redo operations or physical undo operations” based on “comparing the log sequence number to both the first log sequence number and the second log sequence number,” wherein “the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the fist log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the

modified block when the log sequence number is greater than the second log sequence number,” as recited in independent claims 31 and 44.

Weems teaches that a shadow cache is a small separate cache storing address of each value evicted from the main cache. Weems does not teach “performing physical redo operations or physical undo operations” based on “comparing the log sequence number to both the first log sequence number and the second log sequence number,” wherein “the physical redo operations or the physical undo operations are not performed on the modified block when the log sequence number is less than the fist log sequence number, and the physical redo operations are performed on the modified block when the log sequence number is greater than the first log sequence number and less than the second log sequence number, and the physical undo operations are performed on the modified block when the log sequence number is greater than the second log sequence number,” as recited in independent claims 31 and 44.

Thus, Klein, Ganesh, Natarajan, and Weems do not teach, suggest, or render obvious all of the claimed limitations of independent claims 31 and 44.

The examiner also relied on Hayashi et al., U.S. 5,715,447 (“Hayashi”), “*The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition*,” IEEE Press, 2000 (“IEEE”), Raz, U.S. 5,701,480 (“Raz”), and DeWitt, Jr. et al., U.S. 7,093,081 (“DeWitt”) to reject various dependent claims and a portion of independent claim 44. None of these references teach or render obvious the allowable claim features.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue

fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

***Conclusion***

Any inquiry concerning this communication or earlier communications from the Examiner should be directed to AARON SANDERS whose telephone number is (571)270-1016. The Examiner can normally be reached on M-F 9:00a-5:00p.

If attempts to reach the Examiner by telephone are unsuccessful, the Examiner's supervisor, Tim Vo can be reached on 571-272-3642. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Tim T. Vo/  
Supervisory Patent Examiner, Art Unit  
2168

/Aaron Sanders/  
Examiner, Art Unit 2168  
29 April 2010

Application/Control Number: 10/710,356  
Art Unit: 2168

Page 11